On Consensus and Exponentially Fast Social Learning

Pooya Molavi, Kamiar Rahnama Rad, Alireza Tahbaz-Salehi, and Ali Jadbabaie

Abstract— **We analyze a model of social learning in which agents desire to identify an unknown state of the world using both their private observations and information they obtain when communicating with agents in their social neighborhood. Every agent holds a belief that represents her opinion on how likely it is for each of several possible states to be the true one. At each time period, agents receive private signals, and also observe the beliefs of their neighbors in a social network. They then update their beliefs by integrating the information available to them in a boundedly rational fashion. We show that in spite of agents' making new private observations perpetually and the myopic and local updating rule employed by them, agents will eventually reach consensus in their beliefs. This is proved by first showing that agents' beliefs over any state whose truth is inconsistent with their collective observations go to zero exponentially fast.**

I. Introduction

Individuals form opinions about economic, political, and social issues using both their personal experiences and information they obtain through communication with their friends, neighbors, and colleagues. These beliefs determine the decisions they make when faced with different options. The best course of action available to each agent is, however, oftentimes not obvious and depends on unknown parameters (i.e., "states of the world"). Lack of access to all the relevant information about the unknown variables is a stimulus for individuals to share their opinions in order to learn from personal experiences of others. Agents might not be aware of the quality of the information they obtain when communicating with other agents. This could happen due to lack of knowledge both about the informativeness of other agents' observations and the source of their information. Limited knowledge combined with local and myopic processing of information could result in informed agents being either misled or their learning slowed down, as a result of communicating with uninformed agents.[1] In light of this, an important question one could ask is whether individuals come to hold common beliefs, whether these beliefs are correct, and if so, how quickly beliefs become correct. In this paper we provide answers to these questions.

We base our analysis on a non-Bayesian model of social learning proposed in [3]. We consider a society of n agents trying to learn the true state of the world, which can take values in a finite set. Communications between agents are captured by a weighted and possibly directed network where the weights correspond to the trust agents put in the opinions of their neighbors. Agents have beliefs about the likelihood for each of the possible states of the world being the true one. In the beginning of each time step, agents receive private signals and also observe the beliefs held by their neighbors in the previous time step. However, instead of processing new information in a Bayesian way, agents use a simple rule to update their beliefs: Each agent first forms the Bayesian posterior given her observed private signal, as an intermediate step. She then updates her belief to the convex combination of her Bayesian posterior and the beliefs of her neighbors. In [3] the authors show that—under some assumptions—this update eventually leads to social learning, even in finite networks; namely, agents can eventually forecast the future correctly. Furthermore, they will eventually learn the unknown state, given that it is collectively identifiable.

In spite of not being fully Bayesian, this model provides a tractable framework to study the opinion dynamics for agents who repeatedly receive private signals in addition to observing the opinions of their neighbors. This is our main motivation for considering a non-Bayesian protocol. Repeated Bayesian updating of beliefs in presence of social networks can be—except for very simple networks—computationally complicated to carry out. Part of the complications is because there is no reason to believe that agents know the source of their neighbors' information; rather, Bayesian agents have to form beliefs (and repeatedly update them) over possible networks and agents' signal structures. The complexities of Bayesian updating limit its applicability in practice. This is why the Bayesian social learning literature often focuses on simple networks (see, e.g., [4], [5]) or sequential interactions between agents where each agent updates her belief only once (see, e.g., [6]–[10]).

We complement and strengthen the results of Jadbabaie *et al*. in two different ways while maintaining the same assumptions (which are presented and discussed in Section III).

In Section IV we show that, if the true state of the world is identifiable, the probabilities agents assign to wrong states go to zero exponentially fast. This result signifies that, under some mild assumptions, communications between informed and uninformed agents do not hinder learning for informed agents in a fundamental way: not only agents learn the true state as if they were Bayesian learners with access to all the observations, but also they do so exponentially fast similar to Bayesian agents. Proof of this proposition makes use of results regarding exponential stability of ergodic dynamical systems. We treat the beliefs of agents as the trajectory of a random dynamical system and use a linearization argument to conclude almost sure exponential stability.

In Section V we argue that even if the true state of the world is not identifiable by agents, sharing of opinions will lead agents to reach consensus. This is a direct consequence of exponential convergence of agents' beliefs over states that are distinguishable from the true state: Once agents aggregate all the information that is contained in their observations (which happens exponentially fast), their observations contain no new

This work was supported in parts by ONR MURI No. N000140810747, AFOSR Complex Networks Program, and ARO MURI SWARMS Grant.

P. Molavi and A. Jadbabaie are with the Department of Electrical and Systems Engineering and General Robotics, Automation, Sensing and Perception (GRASP) Laboratory, University of Pennsylvania, Philadelphia, PA 19104-6228, USA {pooya, jadbabaie}@seas.upenn.edu

K. Rahnama Rad is with Department of Statistics, Columbia University, New York, NY 10027, USA kamiar@stat.columbia.edu

A. Tahbaz-Salehi is with Decision, Risk, & Operations Division, Columbia Business School, Columbia University, New York, NY 10027, USA alirezat@columbia.edu

[1]See [1] and [2] for examples of failure of learning and consensus, respectively, as a result of existence of stubborn agents.

information and the Bayes' rule becomes the same as identity map. Thereafter, agents simply average their opinions with those of their neighbors, which would eventually lead them to reach consensus. In Proposition 3 we formalize this intuitive argument.

Our paper is related to the line of research on consensus —both in social and engineering networks—which focuses on whether agents, mobile robots, sensors, etc. that have different initial estimates of an unknown will reach an agreement. Examples include the works by Tsitsiklis [11], Jadbabaie *et al.* [12], Olfati-Saber and Murray [13], and Fagnani and Zampieri [14]. What differentiates our paper is presence of agents that repeatedly use new observations to update their beliefs in a boundedly rational way. This makes the agents' update rule both random and non-linear and necessitates use of tools from the theory of random dynamical systems to analyze the model. This paper is also related to a growing body of literature on learning over social networks, especially those with non-Bayesian updating rules. In this spirit are the works by DeGroot [15], DeMarzo *et al.* [16], Acemoglu *et al.* [1], Golub and Jackson [17], Jadbabaie *et al.* [3], and Rahnama Rad and Tahbaz-Salehi [18].

II. MODEL

A. Environment

Our model assumes that time is discrete and there is a finite number of agents, signals, and states of the world. Let Θ be the finite set of possible states of the world. The true state is determined at time zero by nature, and is unchanged thereafter.

Let $\mathcal{N} = \{1, 2, \dots, n\}$ be the set of agents. At time $t \geq 0$ each agent i has a belief about the true state, denoted by $\mu_{i,t}(\theta)$, which is a probability distribution over Θ.

At each time period every agent observes a private random signal. We let $\omega_{i,t} \in S_i$ denote the private signal observed by agent i at time t, where S_i is the set of possible signals for agent i. We let S denote $S_1 \times S_2 \times \cdots \times S_n$. Conditional on $\theta \in \Theta$ being the state of the world, the observation profile $\omega_t = (\omega_{1,t}, \omega_{2,t}, \dots, \omega_{n,t}) \in S$ is generated according to the likelihood function $\ell(\cdot|\,\theta)$ with $\ell_i(\cdot|\,\theta)$ as its ith marginal. Let $\mathbb{P}^\theta = \ell(\cdot|\,\theta)^{\mathbb{N}}$ be the product measure that determines the realization of sequence of signal profiles conditioned on θ being the state of the world, where $\mathbb{N}$ stands for the set of natural numbers. This definition allows for signals to be correlated among agents at the same time period, but makes them independent over time. We assume that $\ell_i(s_i\,|\theta) > 0$ for all $s_i \in S_i$ and all $\theta \in \Theta$; i.e., likelihood functions have the same support conditioned on any θ being the true state. We make this assumption only to simplify the arguments and our results do not critically depend on it.

The interactions between agents are captured by a directed graph $G = (\mathcal{N}, E)$. Let $\mathcal{N}_i = \{j \in \mathcal{N} : (j, i) \in E\}$ be the set of neighbors of agent i. It is assumed that agent j can observe the belief of agent i if there exists a directed edge from i to j; that is, $(i, j) \in E$.

B. Belief Updates

Each agent i starts with the prior belief $\mu_{i,0}(\theta)$ that θ is the true state of the world. At the end of period t, each agent observes the beliefs of her neighbors. At the beginning of the next period, agent i receives the private signal $\omega_{i,t+1}$, and then uses the following rule to update her belief:

$$\mu_{i,t+1}(\theta) = a_{ii}\mu_{i,t}(\theta)\frac{\ell_i(\omega_{i,t+1}|\,\theta)}{m_{i,t}(\omega_{i,t+1})} + \sum_{j\in\mathcal{N}_i} a_{ij}\mu_{j,t}(\theta), \quad (1)$$

where $m_{i,t}(s_i)$ is defined for any $s_i \in S_i$ as

$$m_{i,t}(s_i) = \sum_{\theta\in\Theta} \ell_i(s_i|\,\theta)\mu_{i,t}(\theta).$$

In the update in (1) each agent updates her belief to a convex combination of her own Bayesian posterior given only her private signal and neglecting the social network, and her neighbors' beliefs in the previous period. a_{ij} is the weight that agent i assigns to the opinion of agent j, and a_{ii}, called the *self-reliance* of agent i, is the weight she assigns to her Bayesian posterior conditional on her private signal. We assume that $a_{ij} \geq 0$ and $\sum_{j\in\mathcal{N}_i\cup\{i\}} a_{ij} = 1$ for agents' beliefs to remain a probability distribution over Θ after they perform the update.

When there is no arrival of new information, this update rule is simplified to the one used in DeGroot's naïve learning model. Likewise, when $a_{ij} = 0$ for all $j \neq i$, the model is the same as the Bayesian learning model with no network structure.

C. States of the World

Let θ^* be the true state of the world that is determined at time zero by nature. Let $\bar{\Theta}_i = \{\theta \in \Theta : \ell_i(s_i|\,\theta) = \ell_i(s_i|\,\theta^*)$ for all $s_i \in S_i\}$ be the set of states that are *observationally equivalent* to the true state θ^* from the point of view of agent i, and let $\bar{\Theta} = \bar{\Theta}_1 \cap \cdots \cap \bar{\Theta}_n$ be the set of states that are observationally equivalent to the true state of the world from the point of view of all agents.

We use $\bar{\mu}_{i,t}$ to denote the restriction of $\mu_{i,t}$ to the states which are observationally distinguishable from the true state by at least one of agents. More explicitly, let $\bar{\mu}_{i,t}(\theta)$ be a sub-probability measure over $\Theta \setminus \bar{\Theta}$ such that $\bar{\mu}_{i,t}(\theta) = \mu_{i,t}(\theta)$, for all $\theta \in \Theta \setminus \bar{\Theta}$. Whenever there is no risk of confusion, we refer to both $\mu_{i,t}$ and $\bar{\mu}_{i,t}$ as the belief of agent i at time t.

The evolution of $\bar{\mu}_{i,t}$ is independent of the beliefs of agents on the states that are not distinguishable from θ^*; that is, knowledge of $\bar{\mu}_{j,t}$, for all $j \in \mathcal{N}$, in addition to $\omega_{i,t}$, is sufficient to uniquely determine $\bar{\mu}_{i,t+1}$. To see this note that $m_{i,t}(s_i)$ can be written only in terms of $\bar{\mu}_{i,t}$ as

$$m_{i,t}(s_i) = \ell_i(s_i|\theta^*)\left(1 - \sum_{\theta\in\Theta\setminus\bar{\Theta}} \bar{\mu}_{i,t}(\theta)\right) + \sum_{\theta\in\Theta\setminus\bar{\Theta}} \ell_i(s_i|\theta)\bar{\mu}_{i,t}(\theta). \quad (2)$$

Therefore, by (1) and (2), $\bar{\mu}_{i,t+1}(\theta)$ can be written *only* in terms of $\{\bar{\mu}_{j,t}(\theta)\}_{j\in\mathcal{N}_i}$ and $\omega_{i,t}$.

D. Some Notation

$(\Omega, \mathcal{F}, \mathbb{P})$ is the probability triple, where $\Omega = (\prod_{i=1}^n S_i)^{\mathbb{N}}$, $\mathcal{F}$ is the smallest σ-field that makes all $\omega_{i,t}$ measurable, and $\mathbb{P} = \mathbb{P}^{\theta^*}$ is the probability distribution determining the realization of signals. Let $\mathbb{E}$ be the expectation operator corresponding to $\mathbb{P}$. We use $\omega \in \Omega$ to denote the infinite signal sequence $(\omega_1, \omega_2, \dots)$. Let $\mathcal{F}_{i,t} = \sigma(\omega_1, \omega_2, \dots, \omega_t)$ be the filtration generated by the observations of agents up to time t.

It is sometimes more convenient to use vector notation. We use A to denote the $n \times n$ matrix with the element in ith row and jth column given by a_{ij}, and use $\mu_t(\theta)$ to denote the n dimensional column vector whose ith element is $\mu_{i,t}(\theta)$. We also use μ_t to denote the $n|\Theta|$ dimensional column vector obtained by concatenating the vectors representing $\mu_t(\theta)$ for different $\theta \in \Theta$. Likewise, $\bar{\mu}_t(\theta)$ and $\bar{\mu}_t$ are n and $n|\Theta \setminus \bar{\Theta}|$ dimensional column vectors, respectively, defined in a similar ways. Furthermore, we use $\mathbf{1}$ to denote the column vector with all elements equal to one, and the superscript T to denote the transpose of a vector. Unless otherwise specified, $\|\cdot\|$ denotes an arbitrary vector norm as well as the corresponding induced norm on linear operators.[2] Equation (1) can be written in vector form as

$$\mu_{t+1}(\theta) = A\mu_t(\theta) + \text{diag}\left(a_{ii}\left(\frac{\ell_i(\omega_{i,t+1}|\,\theta)}{m_{i,t}(\omega_{i,t+1})} - 1\right)\right)\mu_t(\theta), \tag{3}$$

where $\text{diag}(v_i)$ is the diagonal matrix with ith diagonal element equal given by v_i.

Given any fixed $\omega \in \Omega$, the belief sequence $\bar{\mu}_0, \bar{\mu}_1, \ldots$ forms the trajectory of the dynamical system $\varphi(\omega) : \mathbb{N} \times \mathbb{R}^{n|\Theta \setminus \bar{\Theta}|} \mapsto \mathbb{R}^{n|\Theta \setminus \bar{\Theta}|}$, where $\varphi_t(\omega)$ is the function that maps the time $t-1$ belief $\bar{\mu}_{t-1}$ to the time t belief $\bar{\mu}_t$. Interpreting agents' beliefs as the trajectory of a dynamical system enables us to use tools from dynamical systems theory to analyze the evolution of the beliefs. This is the approach we take while studying the rate of learning.

III. Assumptions

We maintain the following assumptions throughout the paper.

Assumption 1: The social network is *strongly connected*.[3]

This assumption allows for information to flow from any agent to any other one. One can always assume connectivity without loss of generality, since otherwise each connected component could be analyzed separately. Strong connectivity, on the other hand, requires that any agent that influences other agents be influenced back by them, either directly or indirectly. It excludes the scenarios where some agents place zero total weight on the beliefs of all other agents in the network. If there are agents who influence without being influenced back, then the society can be misled. A simple example of such outcomes is studied in [3].

Assumption 2: there exists at least one agent with positive prior belief on the true state θ^*.

This assumption is similar to the condition that is commonly known as a "grain of truth" in agents' prior beliefs which requires all agents to have positive priors over θ^*. If Assumption 2 is violated, we have the rather uninteresting outcome where all agents continue to have zero belief in the true state at all time periods.

Assumption 3: Self-reliance of all agents is strictly positive.

The purpose of this assumption is twofold. First, the requirement for all agents to have positive self-reliance is sufficient to ensure that the network is aperiodic.[4] which is a necessary condition for agents to reach consensus in their beliefs. Second, Assumption 3 requires agents not to disregard their private observations. If all agents fail to incorporate their observations into their beliefs, information agents cannot accumulate the information that is needed for them to learn the unknown state.

One can show that these assumptions can be weakened significantly by relaxing the strong connectivity assumption and requiring only informed agents to have positive self-reliance [19]. There only needs to be "sufficient" flow of information in the social network; however, to simplify the proofs we maintain Assumptions 1–3 in this paper.

IV. Learning

A central question of interest in study of learning in social networks is whether agents using a certain update rule will learn the true state of the world, and if they learn the state, how fast they do so. Furthermore, one is interested in characterizing the weakest set of assumptions that guarantee learning. In this section we show that if Assumptions 1–3 are satisfied, agents learn an identifiable state *asymptotically almost surely* and also *exponentially almost surely*.

Several notions of learning have been studied in the literature. Here, we say that agents learn the true state if their beliefs in all states except for the true state θ^* go to zero. [5] In our model, agents who have generic prior beliefs cannot learn the true state in finite time; however, the first result presented here states that agents learn the true state if it is identifiable, asymptotically for almost surely sequences of observations profiles. Our second result, which is the main result of this section, strengthens the first one by showing that agents' beliefs over the incorrect states go to zero exponentially fast.

A. Asymptotic Learning

The following result presents a set of conditions under which beliefs of agents over the states that are distinguishable from the true state converge to zero. The proof together with a thorough discussion can be found in [3].

Proposition 1: If Assumptions 1–3 are satisfied, then for all $i \in \mathcal{N}$,

$$\mu_{i,t}(\theta) \to 0 \quad \text{as} \quad t \to +\infty \qquad \forall i \in \mathcal{N}, \quad \forall \theta \in \Theta \setminus \bar{\Theta},$$

with $\mathbb{P}$-probability one.

We say that the true state of the world is *observationally distinguishable* (or identifiable), if there is no $\theta \in \Theta$ that is observationally equivalent to θ^* from the point of view of all agents; that is,

$$\bar{\Theta} = \{\theta^*\}.$$

Since $\mu_{i,t}(\cdot)$ is a probability distribution over Θ for all i and t, the following corollary of Proposition 1 is immediate.

Corollary 1: If Assumptions 1–3 are satisfied and θ^* is observationally distinguishable, then for all $i \in \mathcal{N}$

$$\mu_{i,t}(\theta^*) \to 1 \quad \text{as} \quad t \to +\infty,$$

with $\mathbb{P}$-probability one.

Note that observational distinguishability is necessary for learning. Agents who are bound to know only the marginal likelihood functions cannot distinguish the states in $\bar{\Theta}$, no matter the update they might use.

[2] If a vector norm on a vector space V is given, then one defines the corresponding induced norm or operator norm on the space of bounded linear operators on V as

$$\|M\| = \sup\{\|Mv\| : v \in V \text{ with } \|v\| = 1\}.$$

[3] A graph is called strongly connected if there exists a directed path from any vertex to any other one.

[4] A network is called aperiodic if the greatest common divisor of the lengths of its simple cycles is one, For more on this see [17].

[5] This is different—and in this model is stronger—than the notion of learning known as "weak merging" of opinions. For a discussion of different notions of learning and how they are related see [20], [21].

This result signifies that for agents in a strongly connected social network to learn an identifiable state of the world, they only need to take their own signals into account in a Bayesian fashion and have non-zero self-reliance (i.e., do not discard their private experiences). The averaging of opinions with their neighbors will eventually lead all agents in the social network to learn the true state as if they had access to everyone's observations.

B. Exponential Learning

It is also important to understand when agents are able to learn the true state "sufficiently fast". For instance, if agents are required to make a decision before an exogenous deadline, it is important for them to be sufficiently sure about the true state by the given time. Likewise, if the state is changing exogenously, efficient learning entails the rate of learning to be much faster than the rate by which the true state varies.

We next show that agents who use update (1) learn an identifiable state exponentially fast; that is, as t goes to infinity, $\mu_{i,t}(\theta)$ goes to zero exponentially fast for all agents and all $\theta \in \Theta \setminus \bar{\Theta}$. The following proposition formalizes this result and provides a characterization of the rate of learning.

Proposition 2: If Assumptions 1–3 are satisfied and $\bar{\mu}_0 \neq 0$, then for all ω in a set of $\mathbb{P}$-probability one and all $\epsilon > 0$,

$$\limsup_{t \to +\infty} \frac{1}{t} \log \|\bar{\mu}_t\| \leq \lambda_1 + \epsilon,$$

where $\lambda_1 < 0$ is the (deterministic) top Lyapunov exponent of $M_t(\omega) = D\,\varphi_t(\omega; 0)$, the Jacobian matrix of $\varphi_t(\omega)$ evaluated at the origin.

We prove the proposition by first proving two auxiliary lemmas. Note that the origin is a fixed point of $\varphi(\omega)$ for all $\omega \in \Omega$. Moreover, for any $\omega \in \Omega$, the dynamical system $\varphi(\omega)$ can be decomposed into the sum of its linear part and a higher order term as

$$\varphi_t(\omega; x) = M_t(\omega)\, x + f_t(\omega; x), \tag{4}$$

where $M_t(\omega)$ is the Jacobian of $\varphi_t(\omega; x)$ evaluated at the origin. We first show that the linear dynamical system $M(\omega)$ is asymptotically stable.

Lemma 1: If Assumptions 1 and 3 are satisfied, then for $\mathbb{P}$-almost all ω the linear dynamical system $M(\omega)$ is globally asymptotically stable.

Proof: Let $\Phi(\omega; t)$ be the state transition matrix corresponding to the linear dynamical system $M(\omega)$ defined as

$$\Phi(\omega; t) = M_t(\omega) \dots M_1(\omega).$$

Furthermore, for any $\omega \in \Omega$ let $z_0, z_1, \dots$ denote the trajectory of $M(\omega)$ with initial state z_0. We suppress the dependence of z_t on ω in the rest of the proof. The linear map $M_t(\omega) = D\,\varphi_t(\omega; 0)$ can be represented as a block diagonal matrix with diagonals consisting of $n \times n$ matrices

$$M_t^{\theta}(\omega) = A + \operatorname{diag}\left(a_{ii}\left(\frac{\ell_i(\omega_{i,t}|\,\theta)}{\ell_i(\omega_{i,t}|\,\theta^*)} - 1\right)\right).$$

Therefore, z_t can be written in terms of z_{t-1} as

$$z_t(\theta) = A z_{t-1}(\theta) + \operatorname{diag}\left(a_{ii}\left(\frac{\ell_i(\omega_{i,t}|\,\theta)}{\ell_i(\omega_{i,t}|\,\theta^*)} - 1\right)\right) z_{t-1}(\theta). \tag{5}$$

To prove almost sure global asymptotic stability, we have to show that z_t goes to zero on a set of $\mathbb{P}$-measure one and for *all* initial states $z_0 \in \mathbb{R}^{n|\Theta \setminus \bar{\Theta}|}$. However, we initially assume that $z_{i,0}(\theta)$ is non-negative for all $\theta \in \Theta \setminus \bar{\Theta}$ and $i \in \mathcal{N}$. For all $t \in \mathbb{N}$ and $\omega \in \Omega$, positive orthant is invariant under the action of $\Phi(\omega; t)$. Consequently, z_t is non-negative for all $t \in \mathbb{N}$. This enables us to use the martingale convergence theorem.

A is a stochastic matrix. Moreover, since the network is strongly connected, A is irreducible.[6] Hence, A has a positive left eigenvector v whose corresponding eigenvalue is equal to one.[7] Left multiplying both sides of (5) by v,

$$v^T z_t(\theta) = v^T z_{t-1}(\theta) + \sum_{i=1}^{n} v_i a_{ii}\left(\frac{\ell_i(\omega_{i,t}|\,\theta)}{\ell_i(\omega_{i,t}|\,\theta^*)} - 1\right) z_{i,t-1}(\theta).$$

Taking conditional expectations of the above equation results in

$$\mathbb{E}\left[v^T z_t(\theta) | \mathcal{F}_{t-1}\right] = v^T z_{t-1}(\theta),$$

which since $v^T z_t(\theta) \geq 0$ implies that

$$\mathbb{E}\left[|v^T z_t(\theta)|\right] = \mathbb{E}\left[v^T z_t(\theta)\right] = v^T z_0(\theta) < \infty.$$

Therefore, $v^T z_t(\theta)$ is a martingale with respect to the filtration $\mathcal{F}_{t-1}$. It is also non-negative since v and $z_t(\theta)$ are positive as a result of Perron-Frobenius theorem and our assumption above, respectively. Hence, by martingale convergence theorem, $v^T z_t(\theta)$ converges $\mathbb{P}$-almost surely. Consequently, the martingale difference converges to zero $\mathbb{P}$-almost surely; i.e.,

$$\sum_{i=1}^{n} v_i a_{ii}\left(\frac{\ell_i(\omega_{i,t}|\,\theta)}{\ell_i(\omega_{i,t}|\,\theta^*)} - 1\right) z_{i,t-1}(\theta) \to 0 \quad \text{as} \quad t \to +\infty.$$

The dominated convergence theorem for conditional expectations implies that as $t \to +\infty$,

$$\mathbb{E}\left[\left|\sum_{i=1}^{n} v_i a_{ii}\left(\frac{\ell_i(\omega_{i,t}|\,\theta)}{\ell_i(\omega_{i,t}|\,\theta^*)} - 1\right) z_{i,t-1}(\theta)\right| |\mathcal{F}_{t-1}\right] \longrightarrow 0,$$

$\mathbb{P}$-almost surely, which can be written as

$$\sum_{s \in S} \ell(s|\,\theta^*) \left|\sum_{i=1}^{n} v_i a_{ii}\left(\frac{\ell_i(s_i|\,\theta)}{\ell_i(s_i|\,\theta^*)} - 1\right) z_{i,t-1}(\theta)\right| \longrightarrow 0,$$

$\mathbb{P}$-almost surely. Therefore, since $\ell_i(\cdot|\cdot)$ is strictly positive, for all $(s_1, \dots, s_n) \in S_1 \times \dots \times S_n$ with $\mathbb{P}$-probability one,

$$\sum_{i=1}^{n} v_i a_{ii}\left(\frac{\ell_i(s_i|\,\theta)}{\ell_i(s_i|\,\theta^*)} - 1\right) z_{i,t-1}(\theta) \longrightarrow 0 \quad \text{as} \quad t \to +\infty.$$

Now choose $\hat{s}_i = \arg\max_{s_i \in S_i} \ell_i(s_i|\,\theta)/\ell_i(s_i|\,\theta^*)$. Since $\ell_i(\cdot|\,\theta)$ is a probability distribution over S_i, we have that $\ell_i(\hat{s}_i|\,\theta)/\ell_i(\hat{s}_i|\,\theta^*) \geq 1$ with equality if and only if agent i cannot distinguish θ from θ^*; i.e., $\theta \notin \bar{\Theta}_i$. Therefore, since a_{ii} and v_i are positive,

$$\left(\frac{\ell_i(\hat{s}_i|\,\theta)}{\ell_i(\hat{s}_i|\,\theta^*)} - 1\right) z_{i,t-1}(\theta) \longrightarrow 0 \quad \text{as} \quad t \to +\infty \quad \forall i \in \mathcal{N},$$

with $\mathbb{P}$-probability one. Consequently, $z_{i,t}(\theta)$ converges to zero $\mathbb{P}$-almost surely, for all i such that $\theta \notin \bar{\Theta}_i$. Since $\theta \in \Theta \setminus \bar{\Theta}$, there exists at least one such agent, call it i. Since $z_{j,t}(\theta)$ is non-negative, (5) implies that $z_{j,t}(\theta)$ also converges to zero $\mathbb{P}$-almost surely for all $j \in \mathcal{N}_i$. Note that this happens

[6] The matrix A is called irreducible if for every pair of indices i and j, there exists a natural number m such that $(A^m)_{ij}$ is not equal to 0.

[7] This is a consequence of Perron-Frobenius theorem. See [22], for instance.

even if $\theta \in \bar{\Theta}_j$. Proceeding inductively and using the strong connectivity assumption implies that $z_{j,t}(\theta)$ converges to zero $\mathbb{P}$-almost surely for all $j \in \mathcal{N}$.

For the proof of the general case, note that z_0 can be decomposed as $z_0 = z_0^+ - z_0^-$, where z_0^+ and z_0^- are non-negative. Since $\Phi(\omega; t)$ is a linear map, z_t is given by

$$z_t = z_t^+ - z_t^- = \Phi(\omega; t) z_0^+ - \Phi(\omega; t) z_0^-,$$

which by the above result goes to zero on a set of $\mathbb{P}$-probability one as $t \to +\infty$. ■

The next lemma shows that $f_t(\omega; x) = \varphi_t(\omega; x) - M_t(\omega)x$, as defined in (4), is negligible when x is small. For any $s \in S$ and $\theta \in \Theta$, let $G^\theta(s) : \mathbb{R}^{n|\Theta \setminus \bar{\Theta}|} \mapsto \mathbb{R}^{n \times n}$ be a diagonal matrix valued function whose ith diagonal element is given by

$$G_i^\theta(s; x) = a_{ii} \left(\frac{\ell_i(s_i|\,\theta)}{g_i(s_i; x_i)} - \frac{\ell_i(s_i|\,\theta)}{\ell_i(s_i|\,\theta^*)} \right), \tag{6}$$

where $g_i(s_i; x_i)$ is given by

$$g_i(s_i; x_i) = \ell_i(s_i|\,\theta^*) \left(1 - \sum_{\tilde{\theta} \in \Theta \setminus \bar{\Theta}} x_i(\tilde{\theta}) \right) + \sum_{\tilde{\theta} \in \Theta \setminus \bar{\Theta}} \ell_i(s_i|\,\tilde{\theta}) x_i(\tilde{\theta}).$$

It is straightforward to show that $f_t^\theta(\omega; x) = G^\theta(\omega_t; x) x(\theta)$ for all $\theta \in \Theta \setminus \bar{\Theta}$.

Lemma 2: If Assumption 3 is satisfied, then there exists a neighborhood V of the origin and a constant $K^\theta > 0$ such that if $x, y \in V$, then

$$\|G^\theta(s; x) - G^\theta(s; y)\| \leq K^\theta \|x - y\| \qquad \forall s \in S.$$

The proof is standard, and thus, is omitted to save space. We can now conclude that, as a result of asymptotic stability of $M(\omega)$ and $f_t(\omega; x)$ being small, $\varphi(\omega)$ is exponentially stable.

Proof of Proposition 2: First, note that by sub-multiplicativity of norm and triangle inequality

$$\begin{aligned} &\left\| f_t^\theta(\omega; x) - f_t^\theta(\omega; y) \right\| \\ &\quad \leq \left\| G^\theta(\omega_t; y) \right\| \|x(\theta) - y(\theta)\| \\ &\quad\quad + \left\| G^\theta(\omega_t; x) - G^\theta(\omega_t; y) \right\| \|x(\theta)\|. \end{aligned}$$

Since $G^\theta(\omega_t; 0) \equiv 0$ and by Lemma 2,

$$\begin{aligned} &\left\| f_t^\theta(\omega; x) - f_t^\theta(\omega; y) \right\| \\ &\qquad \leq K^\theta \|x(\theta) - y(\theta)\| \left(\|x(\theta)\| + \|y(\theta)\| \right), \end{aligned}$$

which implies that

$$\|f_t(\omega; x) - f_t(\omega; y)\| \leq K \|x - y\| \left(\|x\| + \|y\| \right),$$

for some $K > 0$.

For any ω let $\Phi(\omega; t)$ be the state transition matrix corresponding to $M(\omega)$, and let $\lambda_1(\omega)$ be the top Lyapunov exponent of $\Phi(\omega; t)$ defined as

$$\lambda_1(\omega) = \lim_{t \to +\infty} \frac{1}{t} \log \|\Phi(\omega; t)\|.$$

The Furstenberg-Kesten theorem [23] implies that for $\mathbb{P}$-almost all ω, the above limit exists and is independent of the realization of ω; moreover, by Lemma 1, for $\mathbb{P}$-almost all $\Phi(\omega; t)$ converges to zero. Hence, $\lambda_1(\omega) = \lambda_1 < 0$ on a set of $\mathbb{P}$-probability one.

Therefore, for any $\epsilon > 0$ there exists a neighborhood $V(\omega)$ of the origin and $C(\omega) > 0$ such that for all $\bar{\mu}_{t_0} \in V(\omega)$ and $t \geq t_0$, and $\mathbb{P}$-almost all $\omega \in \Omega$,[8]

$$\|\bar{\mu}_t\| \leq C(\omega) e^{(\lambda_1 + \epsilon)(t - t_0) + 2\epsilon t_0} \|\bar{\mu}_{t_0}\|. \tag{7}$$

Proposition 1 implies that for any neighborhood $V(\omega)$ of the origin, with $\mathbb{P}$-probability one, there exists $T(\omega)$ such that $\bar{\mu}_t \in V$ for all $t \geq T(\omega)$. Therefore,

$$\limsup_{t \to +\infty} \frac{1}{t} \log \|\bar{\mu}_t\| \leq \lambda_1 + \epsilon,$$

with $\mathbb{P}$-probability one. ■

This proposition strengthens Proposition 1 significantly by stating that if the true state of the world is distinguishable, agents will learn it exponentially fast. This result further underscores the strength of the simple update used in enabling social learning: in spite of bounded rationality of agents and the locality of their interactions, all agents in the social network learn the true state as if they were Bayesian agents observing everyone's signals. [9]

V. Consensus

Will agents' beliefs fluctuate forever or will they eventually converge? If agents' beliefs converge, what are the limiting beliefs? Proposition 1 gives a partial answer to these questions by showing that agents' beliefs over states that are observationally distinguishable from θ^* go to zero. It implies that, if $\bar{\Theta} = \{\theta^*\}$, agents will reach consensus exponentially fast with $\mathbb{P}$-probability one. On the other hand if $\bar{\Theta} = \Theta$, agents will reach consensus exponentially fast regardless of their observations. This is since when agents cannot distinguish any of the states from the true state, the update in (1) is the same as the one in DeGroot's model of learning in which if agents have positive self-reliance and the underlying network is strongly connected, agents will reach consensus [17]. Convergence of agents' beliefs to common values in the two extreme cases discussed above suggests that agents always reach consensus. The following proposition shows that this is indeed true.

Proposition 3: If Assumptions 1–3 are satisfied, then for all $\theta \in \Theta$ and $\mathbb{P}$-almost all ω,

$$\mu_{i,t}(\theta) \to \mu_{i,\infty}(\theta) \quad \text{as} \quad t \to +\infty \qquad \forall i \in \mathcal{N},$$

where $\mu_{i,\infty}(\cdot)$ is a probability distribution over Θ. Furthermore, $\mu_{i,\infty}(\theta) = \mu_{j,\infty}(\theta)$, for all $i, j \in \mathcal{N}$ and $\theta \in \Theta$.

Proof: We can assume that $\bar{\mu}_0 \neq 0$; otherwise, agents' update rule is the same as in DeGroot's model, in which case the proposition is true. Moreover, for $\theta \in \Theta \setminus \bar{\Theta}$ the result is true by Proposition 1.

For any $t \in \mathbb{N}$ and $\theta \in \bar{\Theta}$, the belief vector $\mu_t(\theta)$ can be written in terms of agents' earlier beliefs using the matrix valued function G^θ defined in (6) as

$$\begin{aligned} \mu_t(\theta) &= A\mu_{t-1}(\theta) + G^\theta(\omega_t; \bar{\mu}_{t-1}) \mu_{t-1}(\theta) \\ &= \prod_{\tau=1}^{t} \left(A + G^\theta(\omega_\tau; \bar{\mu}_{\tau-1}) \right) \mu_0(\theta). \end{aligned}$$

By Lemma 2, there exists a neighborhood V of the origin and a constant $K^\theta > 0$ such that if $\bar{\mu}_{\tau-1} \in V$, then

$$\left\| G^\theta(\omega_t; \bar{\mu}_{t-1})) \right\| \leq K^\theta \left\| \bar{\mu}_{t-1} \right\|.$$

[8] This is a corollary of Theorems 1 and 2 in [24] and Oseledets' Theorem [25].

[9] The exponent is, however, generally better (and never worse) for a Bayesian observer with access to all the information.

Proposition 1 implies that there exists $T(\omega)$ such that $\bar{\mu}_t \in V$, for all $t > T(\omega)$. Therefore, by Proposition 2, we can conclude that for any $\epsilon > 0$ and for $\mathbb{P}$-almost all ω, there exists $L(\omega)$ such that

$$\left\|G^\theta(\omega_t; \bar{\mu}_{t-1})\right\| \leq L(\omega) e^{(\lambda_1+\epsilon)(t-1)}. \tag{8}$$

Therefore,

$$\sum_{t=1}^{\infty} \left\|G^\theta(\omega_t; \bar{\mu}_{t-1})\right\| < \infty$$

with $\mathbb{P}$-probability one.

One the other hand, since A is an aperiodic irreducible stochastic matrix, $\lim_{t\to+\infty} A^t = \mathbf{1}v^T$, where v is the left eigenvector of A corresponding to the unit eigenvalue [26].

Hence, on a set of $\mathbb{P}$-measure one,

$$\prod_{\tau=1}^{t} \left(A + G^\theta(\omega_\tau; \bar{\mu}_{\tau-1})\right)$$

converges.[10] Consequently, $\mu_t(\theta)$ converges $\mathbb{P}$-almost surely to some $\mu_\infty(\theta)$. Since

$$\mu_\infty(\theta) = \lim_{t\to+\infty} \prod_{\tau=1}^{t} \left(A + G^\theta(\omega_\tau; \bar{\mu}_{\tau-1})\right) \mu_0(\theta),$$

and $A + G^\theta(\omega_t; \bar{\mu}_{t-1})$ converges to A asymptotically $\mathbb{P}$-almost surely, $\mu_\infty(\theta)$ must satisfy

$$\mu_\infty(\theta) = A\mu_\infty(\theta), \tag{9}$$

which implies that $\mu_\infty(\theta)$ is a right eigenvector of A corresponding to the unit eigenvalue. The vector of one is a solution to (9). Since A is an irreducible stochastic matrix, by Perron-Frobenius theorem, any other solution is a multiple of $\mathbf{1}$. Therefore, $\mu_{i,\infty}(\theta) = \mu_{j,\infty}(\theta)$ for all $i, j \in \mathcal{N}$. ■

Even though beliefs of all agents over the states which are observationally distinguishable from θ^* go to zero, for generic prior beliefs, agents' beliefs over the states in $\bar{\Theta}$ do not converge to either zero or one. They are only constrained to be in the convex hull of agents' prior beliefs. Additionally, the consensus beliefs over such states will generally depend both on the distribution of agents' prior beliefs and the realization of their signals. This is unlike the consensus beliefs over the states that are observationally distinguishable from θ^*; as shown in Proposition 1, agents' beliefs over such states converge to zero independently of the realization of agents' signals and their prior beliefs.

VI. Conclusion

We analyzed a model of social learning where agents repeatedly update their beliefs using a simple rule to incorporate new information they obtain through both personal observations and communications with their neighbors. We showed that if the social network is strongly connected, at least one agent has a positive prior on the true state, and agents do not discard their private observations, they will learn the true state if it is observationally distinguishable and reach consensus if it is not. These results signify that under some mild assumptions, agents will reach consensus in their beliefs, in spite of receiving private signals that might not be consistent with the beliefs communicated to them by their neighbors. In order to prove this, we first showed that agents learn to assign zero probabilities to states that are observationally distinguishable from the true state exponentially fast.

An important question that is left open is: how does the speed of learning depends on network topology and agents' signal structure? We defer a detailed analysis of the rate of learning (and also that of consensus) and its dependence on structural properties of the model to a future paper.

[10] This is a consequence of Theorem 6.8 on page 101 of [27].

References

[1] D. Acemoglu, A. Ozdaglar, and A. Parandeh-Gheibi, "Spread of (mis)information in social networks," *Games and Economic Behavior*, vol. 70, no. 2, pp. 194–227, 2010.

[2] D. Acemoglu, G. Como, F. Fagnani, and A. E. Ozdaglar, "Opinion Fluctuations and Disagreement in Social Networks," *SSRN eLibrary*, 2010. [Online]. Available: http://ssrn.com/paper=1682187

[3] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, "'Non-Bayesian Social Learning', Third Version," *SSRN eLibrary*, 2011. [Online]. Available: http://ssrn.com/paper=1916109

[4] D. Gale and S. Kariv, "Bayesian learning in social networks," *Games and Economic Behavior*, vol. 45, no. 2, pp. 329–346, 2003.

[5] S. Choi, D. Gale, and S. Kariv, "Sequential equilibrium in monotone games: A theory-based analysis of experimental data," *Journal of Economic Theory*, vol. 143, no. 1, pp. 302–330, 2008.

[6] A. V. Banerjee, "A simple model of herd behavior," *The Quarterly Journal of Economics*, vol. 107, no. 3, pp. 797–817, 1992.

[7] S. Bikhchandani, D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as informational cascades," *The Journal of Political Economy*, vol. 100, no. 5, pp. 992–1026, 1992.

[8] L. Smith and P. Sørensen, "Pathological outcomes of observational learning," *Econometrica*, vol. 68, no. 2, pp. 371–398, 2000.

[9] A. Banerjee and D. Fudenberg, "Word-of-mouth learning," *Games and Economic Behavior*, vol. 46, no. 1, pp. 1–22, Jan. 2004.

[10] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *The Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, 2011.

[11] J. Tsitsiklis, "Problems in decentralized decision making and computation," Ph.D. dissertation, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1984.

[12] A. Jadbabaie, J. Lin, and A. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Trans. Autom. Control*, vol. 48, no. 6, pp. 988–1001, 2003.

[13] R. Olfati-Saber and R. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Trans. Autom. Control*, vol. 49, no. 9, pp. 1520–1533, 2004.

[14] F. Fagnani and S. Zampieri, "Randomized consensus algorithms over large scale networks," *IEEE J. Sel. Areas Commun.*, vol. 26, no. 4, pp. 634–649, 2008.

[15] M. H. DeGroot, "Reaching a consensus," *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974.

[16] P. M. DeMarzo, D. Vayanos, and J. Zwiebel, "Persuasion bias, social influence, and unidimensional opinions," *The Quarterly Journal of Economics*, vol. 118, no. 3, pp. 909–968, 2003.

[17] B. Golub and M. Jackson, "Naïve learning in social networks and the wisdom of crowds," *American Economic Journal: Microeconomics*, vol. 2, no. 1, pp. 112–149, 2010.

[18] K. Rahnama Rad and A. Tahbaz-Salehi, "Distributed parameter estimation in networks," in *Proceedings of the 49th IEEE Conference on Decision and Control*, 2010, pp. 5050–5055.

[19] P. Molavi and A. Jadbabaie, "Aggregate observational distinguishability is necessary and sufficient for social learning," in *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference*, 2011, pp. 2335–2340.

[20] D. Blackwell and L. Dubins, "Merging of opinions with increasing information," *The Annals of Mathematical Statistics*, vol. 33, no. 3, pp. 882–886, 1962.

[21] E. Lehrer and R. Smorodinsky, "Merging and learning," *Lecture Notes-Monograph Series*, vol. 30, pp. 147–168, 1996.

[22] A. Berman and R. J. Plemmons, *Nonnegative matrices in the mathematical sciences*. New York: Academic Press, 1979.

[23] H. Furstenberg and H. Kesten, "Products of random matrices," *The Annals of Mathematical Statistics*, vol. 31, no. 2, pp. 457–469, 1960.

[24] L. Barreira and C. Valls, "Stability theory and Lyapunov regularity," *Journal of Differential Equations*, vol. 232, no. 2, pp. 675–701, 2007.

[25] V. Oseledets, "Oseledets theorem," *Scholarpedia*, vol. 3, no. 1, p. 1846, 2008.

[26] C. Meyer, *Matrix analysis and applied linear algebra: solutions manual*. Society for Industrial and Applied Mathematics, 2000, vol. 2.

[27] D. J. Hartfiel, *Nonhomogeneous matrix products*. World Scientific, 2002.